



AFRL-AFOSR-JP-TR-2017-0045

Budgeted Interactive Learning

Hsuan-Tien Lin
NATIONAL TAIWAN UNIVERSITY

06/15/2017
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ IOA
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Executive Services, Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</p>				
1. REPORT DATE (DD-MM-YYYY) 16-06-2017		2. REPORT TYPE Final	3. DATES COVERED (From - To) 21 May 2015 to 20 May 2017	
4. TITLE AND SUBTITLE Budgeted Interactive Learning		5a. CONTRACT NUMBER 5b. GRANT NUMBER FA2386-15-1-4012 5c. PROGRAM ELEMENT NUMBER 61102F		
6. AUTHOR(S) Hsuan-Tien Lin		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NATIONAL TAIWAN UNIVERSITY 1, ROOSEVELT RD., SEC. 4 TAIPEI CITY, 10617 TW			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AOARD UNIT 45002 APO AP 96338-5002			10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR IOA 11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-JP-TR-2017-0045	
12. DISTRIBUTION/AVAILABILITY STATEMENT A DISTRIBUTION UNLIMITED: PB Public Release				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT <p>This project had two questions it was seeking to answer: how to enrich the protocols for interactive learning?, and how to properly make multi-criteria decisions during the interactive learning process? Towards answering the first question, the PI's team broke it into three sub-areas: (1) protocols that combine the benefits of online and batch learning, (2) protocols that improve interactive learning with other sources of information, and (3) protocols that allow extracting useful representations during interactive learning. Aligned with the three sub-areas, they have designed algorithms that allow selecting active learning approaches on the fly (for 2) and transferring the selection experience to other active learning tasks (for 1, 2, and 3). The selection scheme is implemented and released as an open-source active learning package. They have studied theories for designing algorithms for interactive learning with batch-like feedback (for 1) and algorithms for online digestion of representation (for 1 and 3). The team has also addressed real-world needs for considering concept drift during online learning (for 2) and utilizing costs during deep learning, multi-label learning and active learning (for 2 and 3). For the second question, the PI's team has seen promising results on (4) the annotation-budget-sensitive active learning, (5) rethinking deep learning models that trade training/prediction time with performance in large-scale learning, and (6) label embedding models that trade time (embedding length) with performance.</p>				
15. SUBJECT TERMS <p>Machine Learning</p>				
16. SECURITY CLASSIFICATION OF: a. REPORT Unclassified		b. ABSTRACT Unclassified	c. THIS PAGE Unclassified	17. LIMITATION OF ABSTRACT SAR 18. NUMBER OF PAGES 12 19a. NAME OF RESPONSIBLE PERSON KNOPP, JEREMY 19b. TELEPHONE NUMBER (Include area code) 315-227-7006

"Budgeted Interactive Learning"

May/19/2016

Name of Principal Investigators (PI and Co-PIs): Hsuan-Tien Lin

- e-mail address : htlin@csie.ntu.edu.tw
- Institution : National Taiwan University
- Mailing Address : Room 314, CSIE Building, #1 Roosevelt Rd. Sec. 4, Taipei 106, Taiwan
- Phone : +886-2-3366-4888 ext. 314
- Fax : +886-2-3366-4898

Period of Performance: 21-05-2015 – 20-05-2017

Abstract: In this project, we focus on two goals: "how to enrich the protocols for interactive learning?", and "how to properly make multi-criteria decisions during the interactive learning process?" We have richer results on the first goal, which includes three sub-tasks: (1) protocols that combine the benefits of online and batch learning, (2) protocols that improve interactive learning with other sources of information, and (3) protocols that allow extracting useful representations during interactive learning. Aligned with the three sub-tasks, we have designed algorithms that allow selecting active learning approaches on the fly (for 2) and transferring the selection experience to other active learning tasks (for 123). The selection scheme is implemented and released as an open-source active learning package. We have studied theories for designing algorithms for interactive learning with batch-like feedback (for 1) and algorithms for online digestion of representation (for 13). We have addressed real-world needs for considering concept drifts during online learning (for 2) and utilizing costs during deep learning, multi-label learning and active learning (for 23). For the second goal, we have started seeing promising results on (4) annotation-budget-sensitive active learning (5) rethink deep learning models that trades training/prediction time with performance in large-scale learning. (6) label embedding models that trades time (embedding length) with performance.

Introduction: Interaction between teachers and students is important for human learning, but the parallel has not been fully established in machine learning. Furthermore, the resource consumption during the learning process is often neglected by learning algorithms. Realistic use of machine learning, however, demands learning algorithms to be active in obtaining data, progressive in digesting information, and cost/budget-sensitive in making decisions. The direction of budgeted interactive learning is drawing pieces of research attention in recent years with its many applications in personalized recommendation and targeted marketing. The project aims on making machine learning more realistic by studying budgeted interactive learning.

Experimental/Theoretical Methodology, Key Results and Discussion:

We briefly separate our discussion to three directions within budgeted interactive learning. The first one is on cost-sensitive learning and we will present our rich results within the direction. The second one is on active learning and we will present our series of work on making active learning more realistic and budget-oriented. The third one is on online learning and we will discuss our diverse works that tackle online learning theoretically and algorithmically.

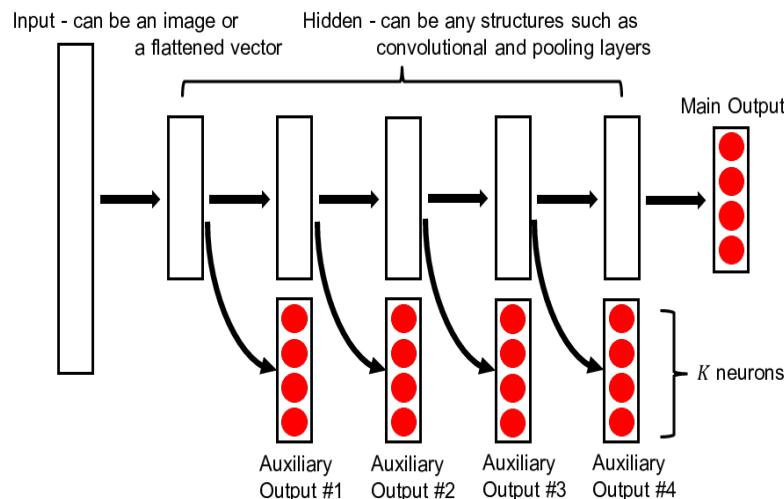
Cost-sensitive learning: We have 8 works related to cost-sensitive learning. One family of

works is on cost-sensitive multi-class classification. In many real-world machine learning applications, classification errors may come with different costs; namely, some types of mis-classification errors may be (much) worse than others. For instance, consider a three-class

classification problem for predicting the state of a patient from {healthy, cold-infected, Zika-infected}. The cost of predicting a Zika-infected patient as healthy shall be remarkably larger

than the cost of predicting a healthy patient as cold-infected, because the former may cause more serious public-health troubles.

Our IJCAI 2016 work [YC2016] advances deep learning towards digesting the cost of mis-classification for batch learning, which relates to utilizing the penalty/reward for different kinds of predictions for interactive learning. Current deep learning models are all cost-insensitive, meaning that they cannot take the cost information into account during training nor prediction. In other words, they cannot distinguish between small mistakes and big mistakes, making it hard to apply them for applications like medical analysis. We take the methodology of designing a novel loss function that effectively reduces cost-sensitive classification to regression. The loss function allows cost-sensitive neural networks to embed the cost information while being sufficiently-smooth for gradient-based optimization. The novel loss function is then plugged into a deep learning model that takes the loss function in both the training and pre-training stages of deep learning. The resulting model is arguably the world's first cost-sensitive deep learning model, and significantly outperforms existing deep learning models as well as alternative cost-sensitive extensions on benchmark cost-sensitive settings. In a sequel work [YC2017] that we have submitted, we further extend the idea and remove the necessity on pre-training. The new idea provides layer-wise cost estimation with auxiliary nodes, and is applicable to a wider range of deep learning architectures, including the convolutional neural network, as illustrated in the following figure.



We have observed promising experimental results based on the layer-wise estimation. For instance, the figure below shows the improvement of the proposed algorithm (with different parameters at the horizontal axis) over the traditional algorithm (the flat line). We see that a 2.5% improvement can be obtained when carefully selecting the parameters in our proposed algorithm. [YC2016] assumes a batch and supervised learning setting where the costs and the labels are fully known, and focus on designing deep learning models that utilize the cost information. We have done other works that tackles the setting when the costs and labels may be unknown. One case is our ECAI work [CY2016], which studies deep reinforcement learning on a special application of bridge bidding. The work can be viewed as a very special

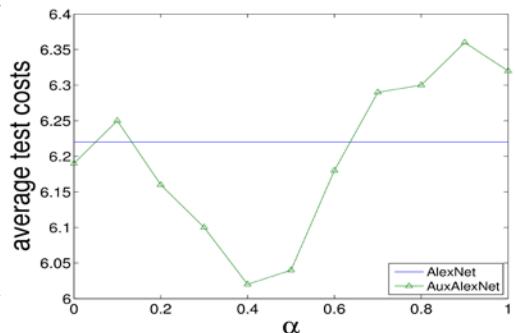
interactive learning system, which mimics how humans interact with each other in practicing mutual understanding. The bridge bidding problem needs a model that learns to be cooperative through exploring the possible costs as the indirect feedback, while exploiting the learned knowledge towards better decision making. We propose a pioneering bridge bidding system without any aid of human domain knowledge. We take the methodology of designing a novel deep reinforcement

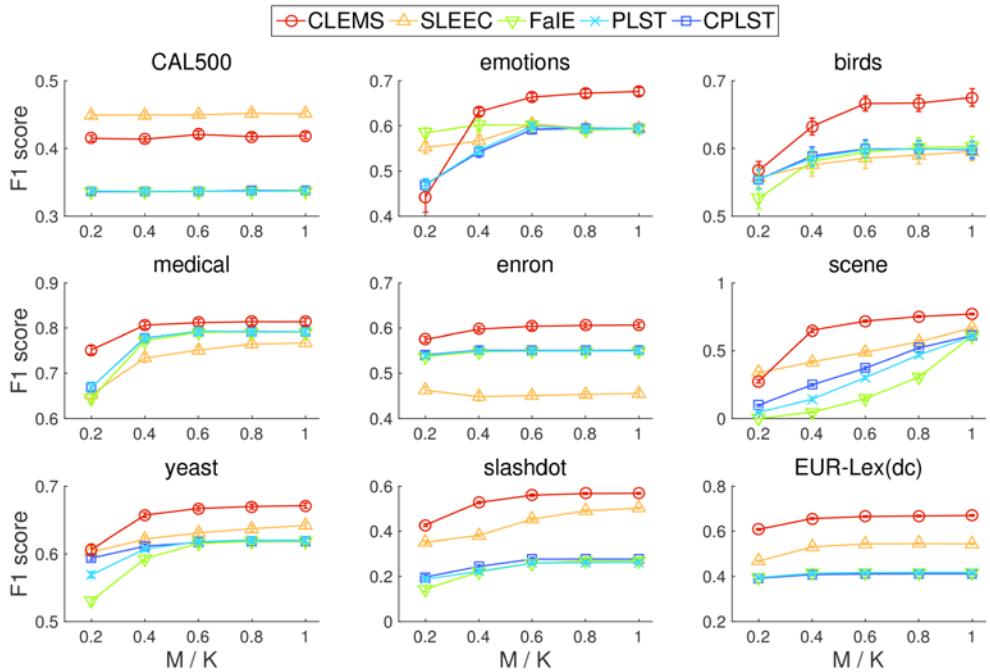
learning model for the system, which extracts sophisticated features and learns to bid automatically based on raw card data. The model includes an upper-confidence-bound algorithm and additional techniques to achieve a balance between exploration and exploitation. Our experiments validate the promising performance of our proposed model. In particular, the model advances from having no knowledge about bidding to achieving superior performance when compared with a champion-winning computer bridge program that implements a human-designed bidding system.

Another case is our ICDM work [KH2016b], which queries the unknown labels---that is, performs active learning, under the cost-sensitive setting. We will illustrate the work in more detail in the active learning direction below.

The works above are on cost-sensitive multi-class classification problems within deep learning, reinforcement learning and active learning. Another family of our cost-sensitive works are on multi-label classification. In particular, by observing that different applications of multi-label classification require different evaluation criteria, we find it important to design general multi-label classification methods that can flexibly take different criteria into account. In our ACML work [YW2016], we propose a novel method that can handle arbitrary example-based evaluation criteria by progressively transforming the cost-sensitive multi-label classification problem into a series of cost-sensitive multi-class classification problems. Experimental results demonstrate that the proposed method is competitive with existing methods under the specific criteria they can optimize, and is superior under several popular criteria.

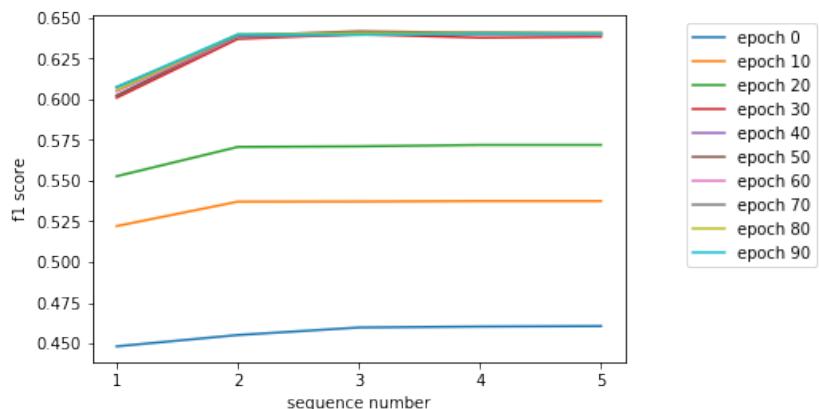
In our recent work [KH2017], we tackle the cost-sensitive multi-label classification problem via another route: label vector embedding. The key idea is similar to [KH2016b], which also embeds the label vectors in some latent space. The proposed algorithm, cost-sensitive label embedding with multidimensional scaling (CLEMS), approximates the cost information with the distances of the embedded vectors by using the classic multidimensional scaling approach for manifold learning. In terms of methodology, CLEMS effectively reduces the cost-sensitive multi-label classification problem to some regression problems. We derive theoretical results that justify how the reduction achieves the desired cost-sensitivity. Furthermore, extensive experimental results demonstrate that CLEMS is significantly better than a wide spectrum of existing LE algorithms and state-of-the-art cost-sensitive algorithms across different cost functions. The nine figures below shows the results that we have got, which represent the trade-off between time (dimension of embedded space in the horizontal axis) and performance (the vertical axis for F1 score). We see that the proposed CLEMS algorithm is much better than other LE algorithms.





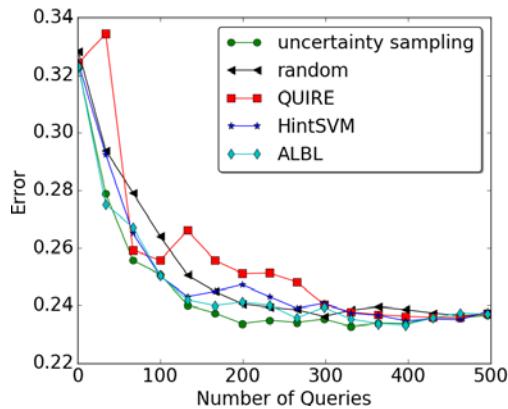
Another recent work of ours [HC2017] moves cost-sensitive multi-label classification and label embedding to the online setting and will be discussed further in the online learning direction below.

One ongoing (unpublished) work of ours is on advancing cost-sensitive multi-label classification models with more sophisticated deep learning techniques. Inspired by the fact that people master some skills for a given set of problem through thinking through the same problem over and over again, we mimic the behavior with a Recurrent Neural Network (RNN) on the multi-label classification problem. In particular, we let the RNN re-think about the previous prediction vector for several times before outputting the final prediction. During the re-thinking process, the costs can be easily fed into the neural network as sample weights to facilitate cost-sensitive learning. Preliminary experimental results in the figure below shows that the cost-sensitive prediction performance (vertical axis) improves after the network rethinks for a few times (horizontal axis). The results can be used to realize budget-sensitive decision making in interactive learning---by playing with the trade-off between the amount of re-thinking and the prediction performance.



Active learning: We have 6 works related to active learning. Active learning allows the learning algorithm to actively query the labels of only a few instances while maintaining good prediction performance. It is a key component of interactive learning that takes human feedback in labeling during the learning process. One practical difficulty of active learning is in selecting proper algorithms/parameters on the fly. In our AAAI work [WH2015], we take the methodology of reducing the online-algorithm-selecting problem as a contextual bandit problem, which is yet another interactive learning problem. We then adopt the EXP4 algorithm with a carefully-designed reward function that calculates a calibrated learning performance of each algorithm to solve the reduced problem. Experimental results demonstrate that the resulting meta-algorithm is often able to select the better algorithms for active learning across different benchmark datasets of active learning.

The outcome of the AAAI work has also lead to libact: an open-source active learning package in Python [YY2017], which has got more than 250 stars on github. The Python package is designed to make active learning easier for general users. The package contains several popular active learning strategies and our AAAI work [WH2015] that assists the users to automatically select the best algorithm/parameter on the fly. Furthermore, the package provides a unified interface for implementing more strategies, models and application-specific labelers. The implementation of our AAAI work (called ALBL) and other active learning algorithms in libact has led to the following demo results, which justify that the ALBL algorithm often matches the best algorithm in terms of active learning performance.

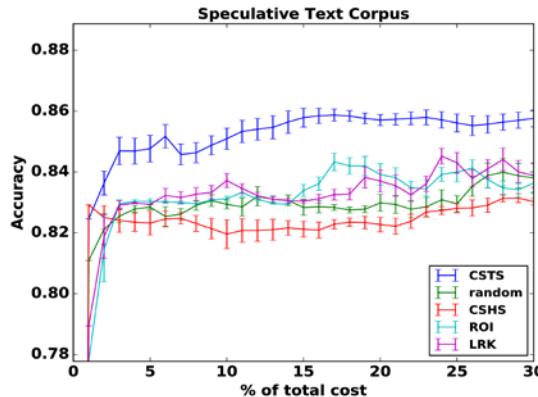


The AAAI work [WH2015] performs randomized selection of algorithms throughout the active learning process of one task only. But the experience of selection cannot be passed to other active learning tasks. In our ICDM work [HC2016], motivated by the philosophical thought that human beings rely on the experience about combining different pieces of knowledge across different active learning tasks, we design algorithms for the machines to do the same. That is, we propose an algorithm that allows the machines to learn a decent combination of different pieces of human knowledge within a single active learning task, and then pass the experience of combination to other tasks to improve the performance of active learning. We take the methodology of reduction again, but this time reducing to a state-of-the-art deterministic algorithm called LinUCB instead of EXP4. We then extend LinUCB to tilt its internal weights towards the experience weights learned from other active learning tasks. The work contributes to the field by proposing a solid definition of what experience means during active learning, and by demonstrating the promising performance of a life-long active learning algorithm across different tasks when compared with state-of-the-art active learning algorithms.

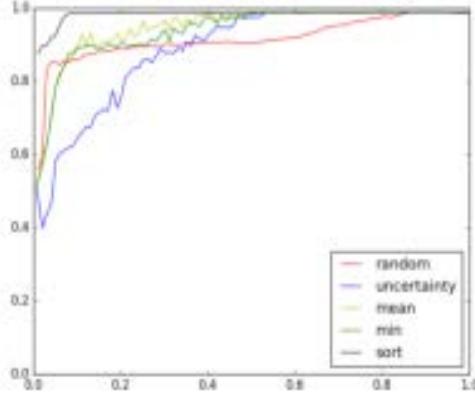
As mentioned, another ICDM work of ours [KH2016b] combines cost-sensitive learning and

active learning. The work can be directly viewed as a working interactive learning system. We propose the world's first non-Bayesian algorithm for tackling the problem. We take the methodology of reducing from cost-sensitive learning to similarity learning by embedding the costs in a latent space via multidimensional scaling, and then calculate the uncertainty of each instance within the latent space. Extensive experimental results demonstrate that the proposed algorithm selects more useful instances by taking the cost information into account through the embedding and is superior to existing cost-sensitive active learning algorithms.

One paper that we have recently submitted [YT2016] is on budgeted active learning. We step out the common assumption that each labeling query is of the same annotation (labeling) cost, and deal with the task where the annotation costs may actually vary between data instances and may be unknown. Traditional active learning algorithms cannot deal with such a realistic scenario. We design a new algorithm that extends the well-known hierarchical sampling algorithm for the task. Our designed algorithm estimates the utility and the cost of each query simultaneously with a tree-structured model motivated from hierarchical sampling. Extensive experimental results over data sets with simulated and true annotation costs validate that the proposed algorithm is generally superior to other annotation-cost-sensitive algorithms. The figure below shows the results on a real-world dataset, where the blue line (CSTS) is the proposed algorithm. We see that the proposed algorithm improves the classification performance (vertical axis) over the budget spent (horizontal axis) much faster than other competitors.



We are in the process of studying another work [SC2017], which focuses on a more general design of annotation-cost-sensitive active learning algorithms. The key methodology is to conduct pre-sampling prior to active learning such that the "expensive" instances can be randomly discarded during pre-sampling. While some results like the figure below demonstrate that some variants of the pre-sampling idea (black line) reaches better performance (vertical axis) over different annotation costs (horizontal axis) than baseline algorithms (red and blue).



Nevertheless, we find that the results are not stable enough for practical use. Thus, we are exploring some ideas on using reinforcement learning to transfer some experience between active learning tasks (similar to what we have done in [HC2016]) to design a more stable version of the pre-sampling algorithm.

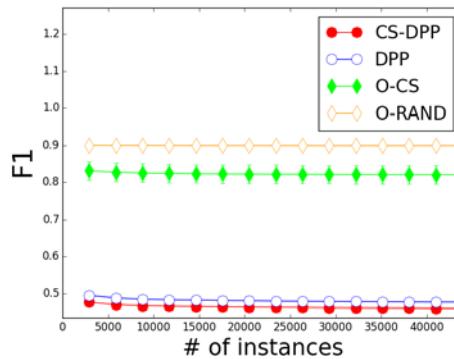
Online learning: We have 4 works related to online learning, which is an important component within interactive learning to digest sequential information. Our online learning works are diverse and range from studying dimension reduction, concept drift, cost-sensitive multi-label learning and bandit decision making.

In our AISTATS work [CL2016], we take a theoretical methodology and study two families of online algorithms to conduct principal component analysis (PCA). Our setup focuses on the memory-restricted setting to be effective for real-world applications. One family updates the PCA matrices in a fully online manner via stochastic gradient descent, the other family updates after a block of sufficient data is gathered, and takes batch PCA algorithm on the block. We advance the first family by generalizing existing theoretical results for arbitrary number of principal components, and advance the second family with designing adaptive block sizes that lead to solid theoretical guarantees. In addition to the theoretical results, we fairly compare the two families and discuss about their cons and pros. The first family enjoys the immediate use of data, and the second family leads to better parameter stability and performance.

In our PAKDD work [SY2016], we address a typical real-world problem of online learning, where the data distribution (concept) can be changing (drifting). There are existing works on detecting the concept drift, but little has been done on what to do after the detection. Other works select more recent data via sliding windows to match the drifting distribution better, but the windows are often fixed regardless of whether the concept drift has been detected or not. The work combines ideas of detection and selection to directly improve the online learning performance under concept drifts. In particular, we take the methodology in designing a meta-algorithm on top of existing online learning algorithms. The novel meta-algorithm un-learns out-dated data to improve the online learning performance, where the un-learning is essentially an automatic mechanism to select proper data. We then extend the un-learning step to design a concept drift detection mechanism by checking the performance difference before and after un-learning. Extensive experimental results demonstrate that the proposed meta-algorithm can be coupled with state-of-the-art online learning algorithms to improve their performance under different kinds of concept drifts.

The previous two works, along with our works on cost-sensitive and multi-label learning with label-space embedding, motivate us to put all the ideas together in a recently submitted work [HC2017]. In this paper, we propose a novel algorithm, cost-sensitive dynamic principal projection (CS-DPP). The algorithm reduces online cost-sensitive classification to online

regression by applying some weighted online PCA on the label space. Particularly, CS-DPP investigates the use of matrix stochastic gradient as the online PCA solver, and establishes its theoretical backbone when coupled with a carefully-designed online regression learner. Practical enhancements of CS-DPP are also studied to improve its effectiveness towards handling the drift of the PCA projection matrix. Experimental results verify that CS-DPP achieves superior practical performance than current MLC algorithms across different evaluation criteria. For instance, the following figure shows that the proposed CS-DPP algorithm is better than the same algorithm without cost-sensitivity (DPP), and is much better than other label space embedding methods across all embedding dimensions (horizontal axis).



Another PAKDD work [KH2016a] of ours stands between online learning and batch learning. In particular, we extend LinUCB, which is a state-of-the-art algorithm for online learning, to a semi-online setting. In a pure online setting, the algorithm is asked to iteratively choose an action based on the observed context, and immediately receives a reward for the chosen action. In real-world applications such as online advertisement, the rewards may not come instantly after choosing an action, and can be received in a pile instead. In this work, we study how LinUCB can be extended for the (semi-online) piled-reward setting to match real-world needs. We contribute to the field by proving the regret bound of a naive use of the original LinUCB algorithm for the piled-reward setting; proposing a novel framework based on the concept of pseudo-rewards to allow more strategic adaptation of the algorithm before the actual rewards come; proving the regret bound of the framework; designing concrete pseudo-rewards for the framework that leads to a novel extension of LinUCB for the piled-reward setting. Experimental results demonstrate that the novel extension leads to significantly better performance on artificial and real-world data sets.

List of Publications and Significant Collaborations that resulted from your AOARD supported project: In standard format showing authors, title, journal, issue, pages, and date, for each category list the following:

a) papers published in peer-reviewed journals,

[YW2017] Yu-Ping Wu and Hsuan-Tien Lin. Progressive k-labelsets for cost-sensitive multi-label classification. Machine Learning 106, no. 5, pages 671-694, May 2017. Accepted via Journal Track of ACML 2016.

b) papers published in peer-reviewed conference proceedings,

[KH2016b] Kuan-Hao Huang and Hsuan-Tien Lin. A novel uncertainty sampling algorithm for cost-sensitive multiclass active learning. In Proceedings of the IEEE International Conference on Data Mining (ICDM), pages 925–930, December 2016.

[HC2016] Hong-Min Chu and Hsuan-Tien Lin. Can active learning experience be

transferred? In Proceedings of the IEEE International Conference on Data Mining (ICDM), pages 841-846, December 2016.

[CY2016] Chih-Kuan Yeh and Hsuan-Tien Lin. Automatic bridge bidding using deep reinforcement learning. In Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI), pages 1362--1369, September 2016.

[YC2016] Yu-An Chung, Hsuan-Tien Lin, and Shao-Wen Yang. Cost-aware pre-training for multiclass cost-sensitive deep learning. In Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI), pages 1411--1417, July 2016.

[CL2016] Chun-Liang Li, Hsuan-Tien Lin, and Chi-Jen Lu. Rivalry of two families of algorithms for memory-restricted streaming pca. In Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS), pages 473--481, June 2016.

[KH2016a] Kuan-Hao Huang and Hsuan-Tien Lin. Linear upper confidence bound algorithm for contextual bandit problem with piled rewards. In Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), volume 1, pages 143--155, April 2016.

[SY2016] Sheng-Chi You and Hsuan-Tien Lin. A simple unlearning framework for online learning under concept drifts. In Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD), volume 1, pages 115--126, April 2016.

[WH2015] Wei-Ning Hsu and Hsuan-Tien Lin. Active learning by learning. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), pages 2659--2665, January 2015.

c) papers published in non-peer-reviewed journals and conference proceedings,

none

d) conference presentations without papers,

none

e) manuscripts submitted but not yet published, and

[SC2017] Si-An Chen and Hsuan-Tien Lin. Annotation-Cost-Sensitive Active Learning by Pre-Sampling, 2017 (working).

[YY2017b] Yao-Yuan Yang, Hong-Min Chu, Yi-An Lin and Hsuan-Tien Lin. Re-think Network for Extreme Cost-Sensitive Multi-label Classification, 2017 (working).

[HC2017] Hong-Min Chu and Hsuan-Tien Lin. Dynamic Principal Projection for Cost-Sensitive Online Multi-label Classification, 2017 (submitted).

[KH2017] Kuan-Hao Huang and Hsuan-Tien Lin. Cost-sensitive Label Embedding for Multi-label Classification, 2017 (submitted).

[YC2017] Yu-An Chung and Hsuan-Tien Lin. Cost-Sensitive Deep Learning with Layer-Wise Cost Estimation, 2017 (submitted).

[YT2017] You-Lin Tsou and Hsuan-Tien Lin. Annotation-Cost-Sensitive Active Learning by Tree Sampling, 2017 (submitted).

[YY2017a] Yao-Yuan Yang, Shao-Chuan Lee, Yu-An Chung, Tung-En Wu, Si-An Chen, and Hsuan-Tien Lin. Libact: Pool-based active learning in python. <https://github.com/ntucllab/libact>, 2017 (submitted).

f) provide a list any interactions with industry or with Air Force Research Laboratory scientists or significant collaborations that resulted from this work.

none

Attachments: Publications a), b) and c) listed above if possible.